

Simultaneous model selection and parameter estimation: A superconducting qubit coupled to a bath of incoherent two-level systems

Markku P.V. Stenberg* and Frank K. Wilhelm

Theoretical Physics, Saarland University, 66123 Saarbrücken, Germany

In characterization of quantum systems, adapting measurement settings based on data while it is collected can generally outperform in efficiency conventional measurements that are carried out independently of data. The existing methods for choosing measurement settings adaptively assume that the model, or the number of unknown parameters, is known. We introduce simultaneous adaptive model selection and parameter estimation. We apply our technique for characterization of a superconducting qubit and a bath of incoherent two-level systems, a leading decoherence mechanism in the state-of-the-art superconducting qubits.

I. INTRODUCTION

Making accurate predictions about nature requires a good model and precise knowledge about its parameters. Model selection [1–3] and parameter estimation are often considered as different goals but Bayesian inference [4] also provides a unified framework where these two tasks become parts of the same question. In this paper, we introduce simultaneous adaptive model selection and parameter estimation. As we explain below, this may be useful, e.g., for characterization of a qubit coupled to a bath of spurious two-level systems, a leading decoherence mechanism in the state-of-the-art superconducting qubits.

In solid state devices one can *a priori* envision many different sources of decoherence and it is important to find systems where quantum information can be protected from decoherence. For a superconducting qubit, a popular early design was based on a Cooper-pair box [5], but it was very susceptible to charge noise, a decoherence mechanism that turned out to be extremely difficult to eliminate. Reengineering the design in a modified Cooper-pair box, a transmon qubit [6], it was, nevertheless, possible to circumvent the problem and considerably enhance the coherence times [7].

In modern transmon designs, experiments have provided considerably specific knowledge on the nature of decoherence. A body of experimental evidence suggests that a leading decoherence mechanism is due to a sparse bath of spurious incoherent two-level systems (TLSs) coupled to a qubit. This evidence consists, e.g., on studies of power [8–10] and frequency [11] dependence of the dielectric loss as well as on the geometrical location of the TLSs [10–14]. Whereas the qubits with a larger junction area, e.g., phase qubits, often couple to TLSs where an excitation may coherently oscillate between the qubit and a TLS [15, 16], such coherent TLSs are rare in transmons [11] where the junctions are smaller. Despite of these advances, the exact microscopic origin of TLSs is still uncertain, and it is important to gather informa-

tion on different aspects [17–20] of individual TLSs. It is therefore useful to be able to efficiently characterize them. Knowing the TLS frequencies precisely also helps to avoid them in controlling the qubit frequency.

Characterization of quantum systems often assumes that the form of the model is known and only its parameters are uncertain. However, there may also exist significant uncertainty even on the form of the model, i.e., on the number of parameters to be estimated. This is the case, for instance, when a qubit is coupled to TLSs whose number is initially unknown. Within Bayesian inference, model selection and parameter estimation may be carried out simultaneously. This fact is built in, e.g., in the recently introduced technique of model averaging [21]. In [21], random, independently on data chosen, measurements were used for state estimation. While data-independent random controls are more easily realized experimentally and form, e.g., the basis of techniques such as randomized benchmarking [22–25], it has been shown theoretically [26–37] and experimentally [38–43] that in many situations adapting measurement settings during data collection can significantly speed up characterization.

For adaptive measurements, the main task is to assign a set of rules, also called a policy, according to which the measurement settings can be chosen efficiently based on the data obtained. For concreteness, we present such a policy for characterization of a qubit and an initially unknown number of TLSs coupled to it. In the Bayesian framework, information on the previous measurements is, rather than revisiting all the previous data, encoded in the current Bayesian posterior probability distribution which can then for adaptive measurements be used for choosing the measurement settings efficiently. The sequential Monte Carlo (SMC) method [44–46] has recently gained popularity in the quantum context [28–34, 36, 37, 42, 43] due to its computational efficiency. But as we explain below, the existing formulation of SMC does not work when the model is initially unknown. We generalize SMC to account simultaneously both for model selection and parameter estimation. Our approach is not limited to a particular physical system but may be used for a larger class of problems.

*markku.stenberg@iki.fi




Figure 1: The background of the figure exhibits the swap spectrum of two incoherent TLSs coupled to a qubit in a numerically simulated experiment with the model parameters randomly chosen within the region described in Sec. VI. The color scale for the swap spectrum represents the qubit excited state probability. The filled circles illustrate a simulated sequence of adaptive measurement settings (cf. Sec. V) in the frequency–waiting time plane. The order of the measurement settings is denoted by their color (from black to white). For clarity, only every tenth measurement setting is shown. The frequencies are measured with respect to a reference frequency ω_0 , cf. Sec. VI. Note the logarithmic scale on the time axis.

II. MODELING INCOHERENT TWO-LEVEL SYSTEMS COUPLED TO A QUBIT

A qubit coupled to a bath of TLSs may be described by a master equation in the Lindblad form [47]

$$\frac{d\hat{\rho}}{dt} = -\frac{i}{\hbar}[\hat{H}_{JC}, \hat{\rho}] + \sum_{j=1}^{n_d} \sum_{k=1}^4 \mathcal{D}[\hat{C}_k^{(j)}]\hat{\rho}, \quad (1)$$

where $\hat{\rho}$ is the density matrix of the qubit and n_d is the number of TLSs coupling to the qubit. When n_d is initially unknown, we follow the terminology in the field of model selection and refer to the model as uncertain, even though the form of Eq. (1) is still assumed to be known. The coherent contribution to the time evolution is given by the Hamiltonian [48]

$$\begin{aligned} \hat{H}_{JC} = & \hbar\omega_q \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) \\ & + \sum_{j=1}^{n_d} \left[\frac{\hbar\omega_d^{(j)}}{2} \hat{\sigma}_z + \hbar g_j \left(\hat{\sigma}_+^{(j)} \hat{a} + \hat{\sigma}_-^{(j)} \hat{a}^\dagger \right) \right], \end{aligned} \quad (2)$$

where $g_j \ll \omega_d^{(j)}$. Here, ω_q is the frequency difference of the ground and the first excited state in the qubit (operators $\hat{a}^{(\dagger)}$) that is assumed to be controllable in the experimental setup and $\omega_d^{(j)}$ is the frequency difference between the two energy levels of the j th TLS ($\hat{\sigma}$ operators). The coupling strength between the qubit and the

j th TLS is denoted by g_j . When the TLSs interact incoherently with the qubit, Markovian decoherence can be described by the Lindblad operator described in Appendix A.

A qubit coupled to incoherent TLSs undergoes frequency-dependent energy relaxation, so that its excited state probability decreases exponentially [11]

$$P_e = e^{-\frac{t}{T_1}}. \quad (3)$$

The directly observable energy relaxation time T_1 of the qubit can be modeled through the frequency-independent energy relaxation time $T_{1,q}$ and the decay rates $\Gamma_{1,d}^{(j)}$ due to different TLSs. For incoherent TLSs, their total contribution is obtained by adding the individual decay rates, and for $\frac{1}{T_{1,q}} < g_j < \frac{1}{T_{2,d}^{(j)}}$ one obtains [11]

$$\frac{1}{T_1} = \frac{1}{T_{1,q}} + \sum_{j=1}^{n_d} \Gamma_{1,d}^{(j)}, \quad \Gamma_{1,d}^{(j)} = \frac{2g_j^2}{1/T_{2,d}^{(j)} + T_{2,d}^{(j)}\Delta_j^2}. \quad (4)$$

Here, $T_{2,d}^{(j)}$ is the coherence time of the j th defect and $\Delta_j = \omega_q - \omega_d^{(j)}$ is the frequency detuning.

III. SWAP SPECTROSCOPY

The understanding of time resolved spectroscopy, referred to as swap spectroscopy [11, 15, 16, 49–51], may be facilitated by a graphical illustration. Figure 1 exhibits a simulated swap spectrum of two incoherent TLSs. One starts by preparing the qubit in the excited state such that the TLSs are in their ground state. This can be done by setting ω_q far from any $\omega_d^{(j)}$ and then exciting the qubit with a microwave pulse. Once ω_q has been fixed to a chosen value (horizontal axis in Fig. 1), the system is allowed to evolve a time t (vertical axis in Fig. 1) after which the qubit is measured in the $\hat{\sigma}_z$ basis. The system is then reset to its ground state before the next measurement. The measurement is repeated at the same setting (ω_q, t) many, usually thousands of times, to approximate the qubit excited state probability (color scale of Fig. 1).

For coherent TLSs, chevron patterns in the swap spectrum can be used to identify the TLS frequency and the strength of its coupling with the qubit [15]. The conventional technique of doing this (for a detailed explanation, see, e.g. [34]) is based on first finding the TLS frequency ω_d and then deducing the coupling strength that is proportional to the frequency of coherent oscillations that the excitation undergoes between the qubit and the TLS. This conventional method does not require Bayesian inference but it is not equally efficient or precise than more sophisticated Bayesian schemes [34, 37]. Furthermore, reviewing the former technique shows that it does not work for incoherent TLSs because an excitation can not coherently oscillate between the qubit and the TLS. Instead, a more general approach based on Bayesian inference can still be applied.

IV. ADAPTIVE BAYESIAN INFERENCE SCHEME

An adaptive Bayesian inference scheme is illustrated in Fig. 2. The starting point is the initial prior probability distribution $P(\mathbf{x}|d_0)$, that describes the experimenters *a priori* conception or subjective belief about the values of the unknown parameters encoded in vector \mathbf{x} and their uncertainties. In principle, assigning the initial prior does not require any data on the current sample, and one may denote $d_0 = \emptyset$. When new data d_{n+1} in the $(n+1)$ th measurement setting is obtained, the updated probability distribution describing the unknown parameters may be obtained through Bayes' theorem [4]

$$P(\mathbf{x}|D_{n+1}) = \frac{P(d_{n+1}|\mathbf{x})P(\mathbf{x}|D_n)}{\int P(d_{n+1}|\mathbf{x}')P(\mathbf{x}'|D_n)d\mathbf{x}'} \quad (5)$$

We assume the measurement is repeated at a single setting M_r times, so that each d_n denotes the outcome of M_r measurement shots. By a measurement shot we mean a single projective measurement with a binary outcome. The ordered set of measurement outcomes in $n-1$ settings is denoted by $D_n = (d_0, \dots, d_n)$. The function $P(d_{n+1}|\mathbf{x})$, referred to as likelihood, describes the probability for data d_{n+1} prior to the experiment for different hypothetical parameter values \mathbf{x} . The likelihood of data $d = e$ (qubit excited state) is one of the key formulas for the adaptive scheme, and in the setup of this paper it is given by Eqs. (3) and (4).

The role of the denominator in Eq. (5) is to ensure the normalization dictated by the conservation of probability. The purpose of applying Bayes' theorem is to obtain the quantity $P(\mathbf{x}|D_{n+1})$, referred to as the posterior, that describes the probability density for different \mathbf{x} given data D_{n+1} . As illustrated in Fig. 2, the posterior may be assigned as the prior before the next measurement which makes iterative application of (5) possible. The estimate is defined by the mean of the posterior $\hat{\mathbf{x}} = \int \mathbf{x}P(\mathbf{x}|D)d\mathbf{x}$. Here, the hat denotes an estimate rather than an operator, and we have omitted the subscript in D for simplicity. Once a sufficient amount of data has been collected, the estimation scheme is terminated and the final estimate is extracted.

V. ADAPTIVE DETECTION AND CHARACTERIZATION OF THE DEFECTS

The Bayesian inference scheme reviewed in the previous section does not, *per se*, dictate how the measurement settings should be chosen. By terminology adopted from machine learning, the rules according to which such a choice is made are referred to as a policy. Note that a policy does not necessarily determine the next measurement setting deterministically but may only give a probability distribution from which the measurement setting is randomly picked. We now present a policy that

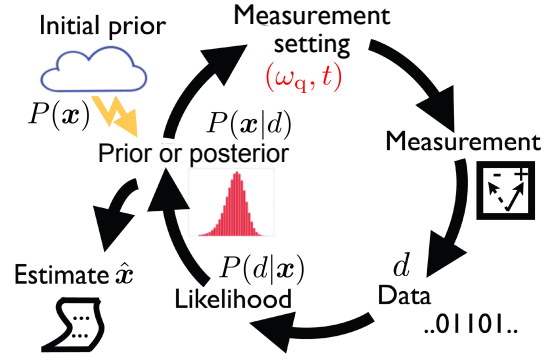


Figure 2: Schematics of an adaptive Bayesian inference scheme with the notation used in the main text.

performs both model selection and parameter estimation simultaneously.

As explained in Sec. II, we allow the number of spurious TLSs to be initially uncertain. Since characterizing different TLSs generally requires measurements with ω_q set to different frequencies, this raises, e.g., a question, how should the measurements be allocated between different potential TLSs. For incoherent TLSs, we also need a different way of choosing the waiting times t than those discussed, e.g., in [34, 37] where it was assumed that an excitation may undergo at least some coherent oscillations between the qubit and the TLS. To quantify the difference between coherent and incoherent TLSs, we note that it can be characterized by the ratio $\frac{g_j}{\Gamma_{1,d}^{(j)}}$. Incoherent TLSs couple incoherently to the qubit and one has $\frac{g_j}{\Gamma_{1,d}^{(j)}} \lesssim \frac{1}{4}$, implying the absence of coherent oscillations between the qubit and the TLS [11].

In the absence of an explicit formula for choosing (ω_q, t) , an experimenter needs to resort to *ad hoc* decisions of his own. To choose (ω_q, t) systematically, we propose the following policy

$$P_{\omega_q} = \frac{P_{1,p}P_{\omega_d^{(1)}}(\omega_q|D) + P_{2,p}P_{\omega_d^{(2)}}(\omega_q|D)}{P_{1,p} + P_{2,p}}, \quad (6)$$

$$t = r\hat{T}_1, \quad r \in (0, 1].$$

Here, ω_q is chosen randomly following the probability distribution P_{ω_q} . The latter is a linear combination of posteriors $P_{\omega_d^{(j)}}$ for the TLS frequencies $\omega_d^{(j)}$. The prefactors $P_{1,p}$ and $P_{2,p}$ denote the prior probabilities for the presence of at least one and two defects, respectively, and the denominator ensures the conservation of probability. Note that if $P_{1,p} + P_{2,p} = 0$, the value of ω_q is irrelevant since by Eqs. (3) and (4), the qubit excited state probability does not depend on frequency. Moreover, r is a uniform random variable on the interval $(0, 1]$ and \hat{T}_1 denotes the mean over the posterior $\hat{T}_1 = \int T_1 P(\mathbf{x}|D)d\mathbf{x}$, with $\mathbf{x} = (g_1, g_2, \omega_d^{(1)}, \omega_d^{(2)}, T_{2,d}^{(1)}, T_{2,d}^{(2)}, T_{1,q})$ encoding all the unknown model parameters. The relaxation time T_1 at different parameter values is obtained through Eq. (4).

In Fig. (1), the filled circles illustrate a simulated sequence of adaptive measurement settings, generated by the policy (6), in the ω_q - t plane. While the settings are initially uniformly distributed on the frequency interval, they later concentrate on the most relevant frequency ranges close to $\omega_d^{(1)}$ and $\omega_d^{(2)}$. This would not be possible with measurement settings chosen prior to data collection which tends to make such measurements less efficient. As directly implied by Eq. (6), the width of the range of waiting times t is proportional to the expected T_1 at a given frequency. Without incorporating the current knowledge about T_1 , one would tend to choose t either too short or too long, rendering the measurement outcomes less sensitive to model parameters and leading to less accurate estimates.

VI. INITIAL PRIOR

To evaluate the performance of our policy, we have applied it to large numbers of simulated samples with the true number of defects n_d equal to 0, 1, and 2. The parameters characterizing a sample are chosen uniformly in random such that $g_1, g_2 \in [0.34, 0.46]$ MHz, $\omega_d^{(1)}, \omega_d^{(2)} \in [\omega_0 - 60, \omega_0 + 60]$ MHz, $T_{2,d}^{(1)}, T_{2,d}^{(2)} \in [50, 100]$ ns, and $T_{1,q} \in [30, 44]$ μ s. Note that the frequency ω_0 denoting the mean of the interval does not play a role in the following discussion. The time interval for $T_{2,d}^{(1)}, T_{2,d}^{(2)}$ is chosen to be consistent with the experimental observations in [9, 11, 16], and the interval for $T_{1,q}$ is chosen to correspond typical values for Xmon qubits [11]. The parameter region for the coupling strengths and the defect frequencies is such that it can contain 2 or a smaller number of defects in experiments with transmons [11].

In the initial prior, the maximum number of defects is assumed to be 2 ($n_d \leq 2$), so that their number can be either 0, 1, or 2, each with the probability $\frac{1}{3}$. Given that the probabilistic information on the different quantities presented in the previous paragraph is everything that is known about a particular sample before data collection, the initial prior is chosen to be consistent with the probability distributions described above which uniquely determines the initial prior. We emphasize that we use the same initial prior for all the figures below, i.e., neither the true number of defects nor the values of any model parameters (except for \hbar) are assumed to be *a priori* known precisely. We label the defects such that the one with the higher (lower) g is referred with the subscript 1 (2). In Figs. 3 and 4, all the curves are obtained by considering 10 000 simulated samples and in Fig. 5 through 100 000 samples.

VII. SEQUENTIAL MONTE CARLO METHOD FOR SIMULTANEOUS ADAPTIVE MODEL SELECTION AND PARAMETER ESTIMATION

We implement the adaptive Bayesian inference scheme numerically through the sequential Monte Carlo (SMC) method [44–46] that has, due to its computational efficiency, recently gained popularity in the context of quantum measurements [28–30, 32–37, 42, 43]. It approximates a continuous probability distribution $P(\mathbf{x}|D)$ through N_p moving grid points or “particles” that are characterized by their locations $\mathcal{S} = \{\mathbf{x}_i\}_{i=1}^{N_p}$ in the parameter space and their relative probabilities, called weights, $\{w_i\}_{i=1}^{N_p}$, that satisfy the normalization condition

$$\sum_{i=1}^{N_p} w_i = 1 \quad (7)$$

due to conservation of probability. We perform the computations with $N_p = 40\,000$. Once new data d_{n+1} has been obtained, Bayes’ theorem (5) implies that the likelihood function may be used to update a weight w_i to its new value w'_i

$$w'_i \sim P(d_{n+1}|\mathbf{x}_n)w_i. \quad (8)$$

To cast this proportionality as an equality, an overall prefactor that is determined by the normalization condition (7) has to be assigned on the right-hand side.

We note that the SMC scheme described, e.g., in [28–30, 32–37, 42–46] assumes a fixed, known, number of parameters to be estimated, and does not therefore adequately apply to our problem where we intend to characterize an *a priori* unknown number of defects in a sample. To point where the problems would arise, let us briefly review the role of the so-called resampling of the particles in the algorithm. Resampling is an element of the SMC scheme but does not constitute the whole algorithm; for the complete presentation of the algorithm, we refer to [29, 37, 52]. Resampling imposes on the particles artificial dynamics whose purpose is to “smoothen” the discrete representation of the probability distribution $P(\mathbf{x}|d_n)$ and to collect the particles in the regions where the probability density is the highest, thus mitigating a limitation in accuracy that a fixed grid would cause. By construction, the artificial dynamics is constrained so that the expected mean and the covariances of the probability distribution are conserved.

Resampling assigns a new particle position in three steps denoted by (i)–(iii) below. These steps are implemented for every particle, i.e., N_p times for each update of the posterior. First, in step (i), among the current particle positions \mathcal{S} , one is chosen randomly following the discrete probability distribution $\{w_i\}_{i=1}^{N_p}$. Let us assume this is the n th choice we make in the current update of the posterior and denote the location by \mathbf{x}_n . In step (ii), the position of this particle is shifted slightly to define

the mean μ_n for the n th resampling distribution

$$\mu_n = a\mathbf{x}_n + (1-a)\hat{\mathbf{x}}. \quad (9)$$

Here, the approximate mean $\hat{\mathbf{x}}$ of the particle locations, i.e., the current estimate of \mathbf{x} , is obtained through $\hat{\mathbf{x}} = \sum_{i=1}^{N_p} w_i \mathbf{x}_i$, and $0 \leq a \leq 1$ is a parameter that determines the location of μ_n on a line connecting \mathbf{x}_n and $\hat{\mathbf{x}}$. For the problem at hand, we obtain the best estimates with approximately $a = 0.995$. The purpose of the second step is to “compress” the probability distribution to counteract the third step that increases the covariances. Finally, in step (iii), the new particle position is assigned by sampling from a (generally multidimensional) normal distribution

$$\mathbf{x}'_n \sim \mathcal{N}[\mu_n, \Sigma] \quad (10)$$

with the mean μ_n and the covariance matrix Σ . The matrix Σ is defined by

$$\Sigma = (1-a^2)\text{Cov}[\mathbf{x}], \quad (11)$$

where $\text{Cov}[\mathbf{x}]$ is the covariance matrix calculated over the particle positions \mathcal{S} . In Eq. (9), also μ_n is a random variable since its value depends on the random variable \mathbf{x}_n , and the role of the prefactor $(1-a^2) \leq 1$ in Eq. (11) is to ensure that the covariances are conserved in resampling.

In our example, we model a sample with two defects ($n_d = 2$) by a vector of the form $\mathbf{x}^{(2)} = (g_1, g_2, \omega_d^{(1)}, \omega_d^{(2)}, T_{2,d}^{(1)}, T_{2,d}^{(2)}, T_{1,q}) \in \mathbb{R}^7$. Since it is computationally convenient that all the particles are represented by vectors of the same length, we represent particles with $n_d = 1$ ($n_d = 0$) by vectors of the form $\mathbf{x}^{(1)} = (g_1, 0, \omega_d^{(1)}, 0, T_{2,d}^{(1)}, 0, T_{1,q})$ [$\mathbf{x}^{(0)} = (0, 0, 0, 0, 0, 0, T_{1,q})$]. This yields the correct observables. We also introduce the following shorthand notation for the sets of vectors describing samples with 0, 1, and 2 defects

$$\begin{aligned} \mathcal{S}^{(0)} &= \{\mathbf{x}_n^{(0)} \in \mathcal{S} | g_{1,n} = 0, g_{2,n} = 0\}, \\ \mathcal{S}^{(1)} &= \{\mathbf{x}_n^{(1)} \in \mathcal{S} | g_{1,n} \neq 0, g_{2,n} = 0\}, \\ \mathcal{S}^{(2)} &= \{\mathbf{x}_n^{(2)} \in \mathcal{S} | g_{1,n} \neq 0, g_{2,n} \neq 0\}. \end{aligned} \quad (12)$$

Since we assume that the number of defects can only have the values $n_d = 0, 1, 2$, we trivially have $\mathcal{S} = \bigcup_{k=0}^2 \mathcal{S}^{(k)}$ and the intersection of the different sets is empty $\mathcal{S}^{(k)} \cap \mathcal{S}^{(j)} = \emptyset$, $k \neq j$.

Let us now consider $\mathbf{x}_n^{(k)}$ ($k = 0, 1$) that initially belongs to $\mathcal{S}^{(0)}$ or $\mathcal{S}^{(1)}$. We note that in resampling, shifting of the vector components in step (ii) through Eq. (9) would generally move $\mathbf{x}_n^{(k)}$ to $\mathcal{S}^{(2)}$. This is clearly an incorrect outcome since steps (ii) and (iii) are merely an unphysical computational procedure. They should not generally immediately change n_d to its maximum value for all the particles, but this is what would happen if the standard formulation of SMC would be applied without any modifications. A logical conclusion is that in step (ii), the components $g_2, \omega_d^{(2)}, T_{2,r}^{(2)}$ (components

$g_1, g_2, \omega_d^{(1)}, \omega_d^{(2)}, T_{2,d}^{(1)}, T_{2,d}^{(2)}$) of $\mathbf{x}_n^{(1)} \in \mathcal{S}^{(1)}$ ($\mathbf{x}_n^{(0)} \in \mathcal{S}^{(0)}$) should not be shifted at all.

But if the SMC algorithm is modified this way, the rest of the scheme has to be modified too, in order to conserve the covariances of the probability distribution. We find that covariances may be conserved by modifying resampling as follows, treating probability distributions for different numbers of particles separately. In step (i) of the resampling, we pick randomly $\mathbf{x}_n \in \mathcal{S}$ as in the standard scheme. As explained above, $\mathbf{x}_n \in \mathcal{S}^{(k)}$ with exactly one k such that we may denote $\mathbf{x}_n = \mathbf{x}_n^{(k)}$. In step (ii), we replace Eq. (9) by the equation $\mu_n^{(k)} = a\mathbf{x}_n^{(k)} + (1-a)\hat{\mathbf{x}}^{(k)}$, where the superscript in $\mu_n^{(k)}$ and $\hat{\mathbf{x}}^{(k)}$ signifies that they are evaluated over $\mathcal{S}^{(k)}$ rather than \mathcal{S} . In (iii), we replace Eq. (10) by $\mathbf{x}'_n \sim \mathcal{N}[\mu_n^{(k)}, \Sigma^{(k)}]$, with $\Sigma^{(k)} = (1-a^2)\text{Cov}[\mathbf{x}^{(k)}]$ where $\text{Cov}[\mathbf{x}^{(k)}]$ is the covariance matrix evaluated over $\mathcal{S}^{(k)}$ rather than \mathcal{S} . When the SMC algorithm is modified this way, steps (ii) and (iii) in resampling conserve the number of defects n_d as well as the expected means and covariances for each $\mathcal{S}^{(k)}$ with $k = 0, 1, 2$ separately. Since the expected means and the covariances are conserved for each $\mathcal{S}^{(k)}$ separately, they are also conserved for \mathcal{S} which satisfies our goal. For the reasons presented above, the computations with an unmodified resampling algorithm would be marred by severe numerical errors, but the modified resampling scheme mitigates these issues. To summarize, one can say that the idea of our generalized SMC scheme is to implement the standard scheme for different sets $\mathcal{S}^{(k)}$ separately.

Within our generalized SMC scheme, the probabilities $P_{k,p}$ ($P_{k,a}$) for the presence of at least k defects (less than k defects) in the sample may be approximated after counting the particles in $\mathcal{S}^{(k)}$, $k = 0, 1, 2$. For instance, let us denote the set of indices of the particles in $\mathcal{S}^{(1)}$ by $\mathcal{I}^{(1)}$ such that $\mathcal{S}^{(1)} = \{\mathbf{x}_i | i \in \mathcal{I}^{(1)}\}$. We can then approximate $P_{1,p} = \sum_{i \in \mathcal{I}^{(1)} \cup \mathcal{I}^{(2)}} w_i$. The probabilities $P_{k,a}$ and $P_{k,p}$ are discussed in the context of Fig. 5.

The computations presented within this article apply the generalized SMC method described above and we carry them out with a computer program using Python and Fortran programming languages and NumPy [53] and SciPy [54] software packages.

VIII. RESULTS

Figure 3 presents the results for samples with exactly two defects ($n_d = 2$). Here, normalized median squared errors are plotted as a function of the number of estimates N_{est} obtained (one for the initial prior and one for each measurement setting). For each measurement setting, the measurement is repeated $M_r = 200$ times, such that the total number of measurement shots is $M = M_r(N_{\text{est}} - 1)$. We find that the errors for $\omega_d^{(1)}$ and $\omega_d^{(2)}$ decrease monotonically over the range of N_{est} studied, but the errors for $g_1, g_2, T_{2,d}^{(1)}, T_{2,d}^{(2)}$ saturate at

N_{est} of the order of $\sim 10^2$. Typical observable relaxation times T_1 of a qubit (an example may be inferred from Fig. 1) are at least an order of magnitude shorter than typical values of $T_{1,q}$ (cf. Sec. VI). Due to relative weakness of the corresponding signal, we are therefore unable to considerably improve the estimates of $T_{1,q}$ from their initial values (not shown in the figure), but the uncertainty in $T_{1,q}$ still allows us to improve the estimates of $\omega_d^{(1)}$ and $\omega_d^{(2)}$ which are experimentally the most relevant quantities.

The samples with a single defect ($n_d = 1$) and with no defects ($n_d = 0$) are described in Fig. 4. A single defect can be characterized by the quantities g_1 , $\omega_d^{(1)}$, and $T_{2,d}^{(1)}$ and the normalized median squared errors of the estimates of these quantities are here plotted as a function of N_{est} . In the absence of any defects, the qubit relaxation is determined by the frequency-independent $T_{1,q}$. At $N_{\text{est}} \sim 10^2$, its median squared error is decreased to the level that with the given initial uncertainty corresponds to the uncertainty of approximately $\sim 1\mu\text{s}$, but with larger N_{est} the error saturates due to issues that we attribute to numerical accuracy. In the presence of a defect, similarly to Fig. 3, due to weakness of the signal, we are not able to considerably improve the estimates of $T_{1,q}$ (not shown in the figure), but the remaining uncertainty in $T_{1,q}$ nevertheless allows to decrease the error of $\omega_d^{(1)}$.

Above, we have not yet shown how quickly the policy learns about the possible absence of defects in the samples. We find that when $M_r = 200$ as above, a single measurement settings is sufficient to find the absence or presence of the defects (not shown in the figures). To more precisely quantify this efficiency, we now reduce the number of repetitions of the measurements, and thereby the information gained, at a single setting and set $M_r = 1$ such that the number of measurement shots is simply $M = M_r(N_{\text{est}} - 1) = N_{\text{est}} - 1$. In Fig. 5, we plot the median Bayesian probabilities for the three different pieces of knowledge describing the absence of defects. These are the probability for the absence of any defects $P_{1,a}$, when the true number of defects is $n_d = 0$ (cyan curve with asterisks), as well as the probabilities for the absence of a second defect $P_{2,a}$, with $n_d = 1$ (brown with '+' symbols) and with $n_d = 0$ (red with crosses). The figure also exhibits median Bayesian probabilities quantifying three different (correct) beliefs in the presence of defects. These are the probability for the existence of at least a single defect $P_{1,p}$, with $n_d = 1$ (blue with circles) and with $n_d = 2$ (black solid), and the probability for the existence of a second defect $P_{2,p}$, when $n_d = 2$ (magenta). The initial values of the probabilities are determined by the initial prior described in Sec. VI. All the probabilities except $P_{2,a}$ with $n_d = 1$ and $P_{2,p}$ with $n_d = 2$ reach a level virtually indistinguishable from unity, i.e., certainty about a belief that is correct, after $M \sim 10$ measurement shots whereas the latter reach such a value after $M \sim 10^2$ shots. Note that here, the distance to unity can be interpreted as error in model selection since it represents the

Bayesian probability of the all the other models except the correct one. The initial prior and Bayes' theorem (5) imply that the possible number of defects allowed in any posterior can not exceed 2. Comparing the results of Fig. 5 with those in Figs. 3 and 4 shows that here, model selection requires considerably less measurements than parameter estimation.




Figure 3: Accuracy of parameter estimation in the presence of two defects in a sample ($n_d = 2$). Normalized median squared error as a function of the number of estimates N_{est} for different parameters of the model described in Sec. II: g_1 (cyan with circles), g_2 (blue with asterisks), $\omega_d^{(1)}$ (red with squares), $\omega_d^{(2)}$ (black solid), $T_{2,d}^{(1)}$ (brown with '+' symbols), and $T_{2,d}^{(2)}$ (magenta with crosses).




Figure 4: Accuracy of parameter estimation in the absence of any defects in a sample ($n_d = 0$) and in the presence of a single defect ($n_d = 1$). Normalized median squared error as a function of the number of estimates N_{est} for different parameters of the model described in Sec. II: g_1 (brown with circles), $\omega_d^{(1)}$ (black solid), $T_{2,d}^{(1)}$ (blue with '+' symbols), and $T_{1,q}$ (red with squares).




Figure 5: Model selection in the presence of an unknown number of defects in a sample, cf. Eq. (1). Different curves exhibit the medians of the Bayesian probabilities as a function of the number of measurement shots. The Bayesian probabilities are for the presence of at least a single defect ($P_{1,p}$), for the presence of a second defect ($P_{2,p}$), for the absence of any defects ($P_{1,a}$), and for the absence of a second defect ($P_{2,a}$). The true number of defects is denoted by $n_d = 0, 1, 2$.

IX. CONCLUSION

We have generalized the concept of adaptive measurements to account for simultaneous model selection and parameter estimation. We implemented a policy built on an adaptive Bayesian scheme using a generalization of sequential Monte Carlo method. Since global optimization through utility [4] maximization would generally be numerically challenging, many recent works making use of Bayesian inference have adopted either greedy algorithms [27, 28, 30, 42, 43, 55–57] that optimize a certain utility function assuming the next measurement is the last one, or different heuristics that choose measurement settings adaptively through a probability distribution [32–34, 36, 37]. Our work belongs to the latter category.

Assuming the model is known and that only its parameter values are uncertain, it should be noted that the difficulty of parameter estimation depends on the model. In certain cases it is possible to drastically outperform nonadaptive measurements and, e.g., to improve the accuracy of the estimate exponentially as a function of the number of measurement shots [27, 30, 32–34]. Basically, this is because in the systems considered in [27, 30, 32–34], small changes in the coupling strength may correspond to a large change in the observable P_e when t is long. However, in the system studied within this paper (with the parameters in the region discussed in Sec. VI), long t leads to loss of the signal due to rapid relaxation of the qubit which is why small changes in g_j and $T_{2,d}^{(j)}$ always correspond to small differences in P_e . Hence inferring g_j and $T_{2,d}^{(j)}$ is difficult. Estimation of $T_{1,q}$ is challenging because the competing contribution of $\Gamma_{1,d}^{(j)}$

(that depends on g_j and $T_{2,d}^{(j)}$) on P_e is much stronger, cf. Eqs. (3) and (4) and Sec. VI. Furthermore, since the volume of the parameter space increases exponentially with its dimension, this makes the set of particles sparser. This “curse of dimensionality” decreases the effective sample size $N_{\text{ess}} = \frac{1}{\sum_i w_i^2}$, making it smaller than, e.g., in the considerations of [34, 36, 37], which may increase the numerical error. We attribute the observed floors in accuracy in Figs. 3 and 4 to competition between information obtained through measurements and accumulated numerical error. Despite these challenges, we are able to decrease the squared error of $\omega_d^{(j)}$ by a factor of $\sim 10^3$ by making only $\sim 10^3$ updates in the measurement setting. We emphasize that $\omega_d^{(j)}$ are the quantities that are the most relevant, e.g., to protect a qubit from unwanted decoherence since avoiding parking the qubit at these frequencies makes its coherence time longer.

Even though in this paper we have considered characterization of incoherent TLSs as an example, the idea of simultaneous adaptive model selection and parameter estimation is applicable to a larger class of problems. For instance, once the required likelihood functions have been derived, it is straightforward to generalize the work in [34, 37] for an unknown number of coherent TLSs.

For a larger number of TLSs in the frequency interval, we note that when their frequencies ω_i are sufficiently far from each others to neglect the overlaps in their swap spectra, a simple approximation holds. Denoting the mean squared error of a TLS frequency ω_i by $\mathcal{E}_{\omega_i}^2$, we expect that $\frac{1}{n_d} \sum_{i=1}^{n_d} \mathcal{E}_{\omega_i}^2(Mn_d)|_{n_d} \approx \mathcal{E}_{\omega_1}^2(M)|_{n_d=1}$, i.e., the mean squared error averaged over n_d defects after Mn_d measurement shots is approximately equal to $\mathcal{E}_{\omega_1}^2$ after M shots in the presence of a single defect because in the former case the measurements have to be allocated among a larger number (n_d) of different TLSs. In a more realistic scenario, one has to take into account the overlaps in the swap spectra which is difficult without doing the full numerical calculation.

Another topic for further research is the influence of imperfections in the readout. Such errors are best tolerated when first quantified and then incorporated in the likelihood function. For instance, when the probability of misidentifying a qubit ground state as an excited state or vice versa equals γ one has to replace Eq. (3) by $P_e = (1 - 2\gamma) \exp[-t/T_1] + \gamma$. As discussed in Sec. VIII, selecting the correct model usually requires much less data than achieving a reasonable degree of accuracy in parameter estimation, i.e., the correct model is usually found in a very early stage of the parameter estimation. On the other hand, adaptive Bayesian inference scheme combined with SMC approximation has been shown to be robust against substantial amount of readout errors for parameter estimation of different fixed models [32–34, 37] as well for estimation of different quantum states [36, 42, 43]. We thus expect that robustness against readout errors can also quite generally be achieved for simul-

taneous adaptive model selection and parameter estimation.

In conclusion, assuming both the model and all its parameters are initially uncertain, we formulated system characterization in a unified Bayesian framework and delivered numerical tools required for solving a concrete problem. We applied our method for efficient detection and characterization of a sparse bath of incoherent spurious two-level systems, a leading decoherence mechanism in the state-of-the-art superconducting qubits.

ACKNOWLEDGEMENTS

We acknowledge O. Köhn for discussions. This work was supported by the European Union through ScaleQIT.

Appendix A: Lindblad operator

When the TLSs interact incoherently with the qubit, Markovian decoherence can be described by the Lindblad operator [58]

$$\mathcal{D}[\hat{C}_k^{(j)}]\hat{\rho} = \hat{C}_k^{(j)}\hat{\rho}\hat{C}_k^{(j)\dagger} - \frac{1}{2}\{\hat{C}_k^{(j)\dagger}\hat{C}_k^{(j)}, \hat{\rho}\}, \quad (\text{A1})$$

where the terms

$$\begin{aligned} \hat{C}_1^{(j)} &= \hat{a}\sqrt{\Gamma_{1,q}^{(j)}}, & \hat{C}_2^{(j)} &= \hat{a}^\dagger\hat{a}\sqrt{2\Gamma_{\phi,q}^{(j)}}, \\ \hat{C}_3^{(j)} &= \hat{\sigma}_-\sqrt{\Gamma_{1,d}^{(j)}}, & \hat{C}_4^{(j)} &= \hat{\sigma}_+\hat{\sigma}_-\sqrt{2\Gamma_{\phi,d}^{(j)}}, \end{aligned} \quad (\text{A2})$$

denote the energy and phase relaxation rates of the qubit and the defect.

-
- [1] K. P. Burnham and D. R. Anderson, *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach* (Springer, New York, 2002).
 - [2] H. Akaike, IEEE Trans. Autom. Control **19**, 716 (1974).
 - [3] G. Schwarz, Ann. Statist. **6**, 461 (1978).
 - [4] D. Sivia and J. Skilling, *Data Analysis: A Bayesian Tutorial* (Oxford University Press, New York, 2006).
 - [5] Y. Nakamura, Yu. A. Pashkin, and J. S. Tsai, Nature (London) **398**, 786 (1999).
 - [6] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, R. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Phys. Rev. A **76**, 042319 (2007).
 - [7] M. H. Devoret and R. J. Schoelkopf, Science **339**, 1169 (2013).
 - [8] M. v. Schickfus and S. Hunklinger, Phys. Lett. A **64**, 144 (1977).
 - [9] J. M. Martinis, K. B. Cooper, R. McDermott, M. Steffen, M. Ansmann, K. D. Osborn, K. Cicak, S. Oh, D. P. Pappas, R. W. Simmonds, et al., Phys. Rev. Lett. **95**, 210503 (2005).
 - [10] R. Barends, N. Vercruyssen, A. Endo, P. J. Visser, T. Zijlstra, T. M. Klapwijk, P. Diener, S. J. C. Yates, and J. J. A. Baselmans, Appl. Phys. Lett. **97**, 023508 (2010).
 - [11] R. Barends *et al.*, Phys. Rev. Lett. **111**, 080502 (2013).
 - [12] J. Gao, M. Daal, A. Vayonakis, S. Kumar, J. Zmuidzinas, B. Sadoulet, B. A. Mazin, P. K. Day, and H. G. Leduc, Appl. Phys. Lett. **92**, 152505 (2008).
 - [13] H. Wang *et al.*, Appl. Phys. Lett. **95**, 233508 (2009).
 - [14] J. Wenner *et al.*, Appl. Phys. Lett. **99**, 113513 (2011).
 - [15] K. B. Cooper, M. Steffen, R. McDermott, R. W. Simmonds, S. Oh, D. A. Hite, D. P. Pappas, and J. M. Martinis, Phys. Rev. Lett. **93**, 180401 (2004).
 - [16] Y. Shalibo, Y. Roife, D. Shwa, F. Zeides, M. Neeley, J. M. Martinis, and N. Katz, Phys. Rev. Lett. **105**, 177001 (2010).
 - [17] T. Palomäki *et al.*, Phys. Rev. B **81**, 144503 (2010).
 - [18] J. Lisenfeld, C. Müller, J. H. Cole, P. Bushev, A. Lukashenko, A. Shnirman, and A. V. Ustinov, Phys. Rev. Lett. **105**, 230504 (2010).
 - [19] G. J. Grabovskij, T. Peichl, J. Lisenfeld, G. Weiss, and A. V. Ustinov, Science **338**, 232 (2012).
 - [20] B. Sarabi, A. N. Ramanayaka, A. L. Burin, F. C. Wellstood, and K. D. Osborn, Phys. Rev. Lett. **116**, 167002 (2016).
 - [21] C. Ferrie, New J. Phys. **16**, 093035 (2014).
 - [22] J. Emerson, R. Alicki, and K. Życzkowski, J. Opt. B **7**, S347 (2005).
 - [23] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Phys. Rev. A **77**, 012307 (2008).
 - [24] E. Magesan, J. M. Gambetta, and J. Emerson, Phys. Rev. Lett. **106**, 180504 (2011).
 - [25] E. Magesan *et al.*, Phys. Rev. Lett. **109**, 080505 (2012).
 - [26] D. W. Berry, B. L. Higgins, S. D. Bartlett, M. W. Mitchell, G. J. Pryde, and H. M. Wiseman, Phys. Rev. A **80**, 052114 (2009).
 - [27] A. Sergeevich, A. Chandran, J. Combes, S. D. Bartlett, and H. M. Wiseman, Phys. Rev. A **84**, 052315 (2011).
 - [28] F. Huszár and N. M. T. Houlshby, Phys. Rev. A **85**, 052120 (2012).
 - [29] C. E. Granade, C. Ferrie, N. Wiebe, and D. G. Cory, New J. Phys. **14**, 103013 (2012).
 - [30] C. Ferrie, C. E. Granade, and D. G. Cory, Quant. Inf. Proc. **12**, 611 (2013).
 - [31] C. Ferrie, Phys. Rev. Lett. **113**, 190404 (2014).
 - [32] N. Wiebe, C. E. Granade, C. Ferrie, and D. G. Cory, Phys. Rev. Lett. **112**, 190501 (2014).
 - [33] N. Wiebe, C. E. Granade, C. Ferrie, and D. G. Cory, Phys. Rev. A **89**, 042314 (2014).
 - [34] M. P. V. Stenberg, Y. R. Sanders, and F. K. Wilhelm, Phys. Rev. Lett. **113**, 210404 (2014).
 - [35] C. E. Granade, C. Ferrie, and D. G. Cory, New J. Phys. **17**, 013042 (2015).
 - [36] M. P. V. Stenberg, K. Pack, and F. K. Wilhelm, Phys. Rev. A **92**, 063852 (2015).
 - [37] M. P. V. Stenberg, O. Köhn, and F. K. Wilhelm, Phys. Rev. A **93**, 012122 (2016).
 - [38] Th. Hannemann, D. Reiss, Ch. Balzer, W. Neuhauser, P. E. Toschek, and Ch. Wunderlich, Phys. Rev. A **65**, 050303(R) (2002).

- [39] B. L. Higgins, D. W. Berry, S. D. Bartlett, H. M. Wiseman, and G. J. Pryde, *Nature* (London) **450**, 393 (2007).
- [40] G. Y. Xiang, B. L. Higgins, D. W. Berry, H. M. Wiseman, and G. J. Pryde, *Nat. Photon.* **5**, 43 (2011).
- [41] H. Yonezawa *et al.*, *Science* **337**, 1514 (2012).
- [42] K. S. Kravtsov, S. S. Straupe, I. V. Radchenko, N. M. T. Houlby, F. Huszár, and S. P. Kulik, *Phys. Rev. A* **87**, 062122 (2013).
- [43] G. I. Struchalin, I. A. Pogorelov, S. S. Straupe, K. S. Kravtsov, I. V. Radchenko, and S. P. Kulik, *Phys. Rev. A* **93**, 012103 (2016).
- [44] M. West, *J. Roy. Stat. Soc. B Met.* **55**, 409 (1993).
- [45] N. J. Gordon, D. J. Salmond, and A. F. M. Smith, *Radar and Signal Processing IEE Proc.-F* **140**, 107 (1993).
- [46] J. Liu and M. West in, *Sequential Monte Carlo Methods in Practice*, edited by A. Doucet, N. Freitas, and N. Gordon (Springer, New York, 2001).
- [47] J. Lindblad, *Commun. Math. Phys.* **48**, 119 (1976).
- [48] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, *Phys. Rev. A* **69**, 062320 (2004).
- [49] M. Neeley, M. Ansmann, R. C. Bialczak, M. Hofheinz, N. Katz, E. Lucero, A. O'Connell, H. Wang, A. N. Cleland, and J. M. Martinis, *Nat. Phys.* **4**, 523 (2008).
- [50] M. Mariantoni *et al.*, *Science* **334**, 61 (2011).
- [51] M. Mariantoni *et al.*, *Nat. Phys.* **7**, 287 (2011).
- [52] J. Liu and M. West in, *Sequential Monte Carlo Methods in Practice*, edited by A. Doucet, N. Freitas, and N. Gordon (Springer, New York, 2001).
- [53] D. Ascher, P. Dubois, K. Hinsien, J. Hugunin, and T. Oliphant, computer code NUMERICAL PYTHON 2001, <http://www.numpy.org>.
- [54] E. Jones *et al.*, computer code SCIPY, 2001, <http://www.scipy.org>.
- [55] D. W. Berry and H. M. Wiseman, *Phys. Rev. Lett.* **85**, 5098 (2000).
- [56] D. G. Fischer and M. Freyberger, *Phys. Lett. A* **273**, 293 (2000).
- [57] D. W. Berry, H. M. Wiseman, and J. K. Breslin, *Phys. Rev. A* **63**, 053804 (2001).
- [58] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, (Cambridge University Press, Cambridge, England, 2000).